

Nested Named Entity Recognition as Corpus Aware Holistic Structure Parsing

Yifei Yang^{1,2,†}, Zuchao Li^{1,2,†}, Hai Zhao^{1,2,*}

¹Department of Computer Science and Engineering, Shanghai Jiao Tong University

²MoE Key Lab of Artificial Intelligence, AI Institute, Shanghai Jiao Tong University
{yifeiyang, charlee}@sjtu.edu.cn, zhaohai@cs.sjtu.edu.cn

Abstract

As a fundamental natural language processing task and one of core knowledge extraction techniques, named entity recognition (NER) is widely used to extract information from texts for downstream tasks. Nested NER is a branch of NER in which the named entities (NEs) are nested with each other. However, most of the previous studies on nested NER usually apply linear structure to model the nested NEs which are actually accommodated in a hierarchical structure. Thus in order to address this mismatch, this work models the full nested NEs in a sentence as a holistic structure, then we propose a holistic structure parsing algorithm to disclose the entire NEs once for all. Besides, there is no research on applying corpus-level information to NER currently. To make up for the loss of this information, we introduce Point-wise Mutual Information (PMI) and other frequency features from corpus-aware statistics for even better performance by holistic modeling from sentence-level to corpus-level. Experiments show that our model yields promising results on widely-used benchmarks which approach or even achieve state-of-the-art. Further empirical studies show that our proposed corpus-aware features can substantially improve NER domain adaptation, which demonstrates the surprising advantage of our proposed corpus-level holistic structure modeling.¹

1 Introduction

Named Entity Recognition is to find predefined named entities such as locations, organizations or people in text, which usually serves as an upstream natural language processing (NLP) task (Huang et al., 2015; Lample et al., 2016; Ma and Hovy,

* Corresponding author. † These authors made equal contribution to this paper. This work was supported by Key Projects of National Natural Science Foundation of China (U1836222 and 61733011).

¹<https://github.com/yangyifei729/NerAsParsing>




Figure 1: An example of nested named entities. The solid lines or brackets represents the start and end positions of a named entity. ‘GPE’ is a named entity category, representing geographic or political entities. *France* and *San Francisco* are nested by *the San Francisco of France* which forms a hierarchy structure.

2016) and one of key knowledge extraction techniques in knowledge engineering.

Early works on NER (Liu et al., 2011; Ritter et al., 2011) mostly cope with only flat NEs that have no overlapped relationship with each other. However, as a common language phenomenon, nested NER appears universally in many corpora such as the field of biology or news events. An instance of overlapped NEs is shown in Figure 1. Since solving the Nested NER task will bring more contextual information than the Flat NER task and promotes downstream tasks, it recently has aroused great research interest (Luo and Zhao, 2020; Muis and Lu, 2017; Wang and Lu, 2018). Although researchers have made good progress on flat NER, existing studies nested NER has not achieved a more satisfactory result yet.

With the development of deep learning technology, nested NER has gradually developed from the methods based on handcraft features and traditional machine learning (Alex et al., 2007; Kumar et al., 2008) to the methods based on neural networks. In recent years, myriad studies apply neural network models to achieve state-of-the-art results.

From a perspective of deep learning, NER may be conveniently cast to a multi-class classification task or sequence labeling task (Fisher and Vlachos, 2019; Lample et al., 2016; Ma and Hovy,

2016). Nevertheless, such modeling ways which mostly suit for flat NER cannot well handle nested NEs. Instead, there are two categories of models for nested NER, Layered-based Model and Region-based Model. The former recognizes the hierarchical structure by stacking multiple Flat NER layers and the latter type enumerates all text spans to discover the possible named entity mentions.

However, all of the above methods still view all nested NEs as a sequence structure and identify all the possible tags one by one for each NE span, which leads to two obvious drawbacks in modeling capability, (1) These methods require multiple modules to cooperate and each module has to be able to recognize NEs from a complete sentence, which usually results in a too complicated model and time-consuming processing (Luo and Zhao, 2020; Wang et al., 2020). (2) When relying on multiple modules working together, these methods further impose substantial information transmission among different modules, which inevitably causes serious error propagation (Ju et al., 2018).

We address above limitation by modeling the nested NEs in a sentence as a holistic structure and propose a model which builds off work from recently constituency parsing algorithm (Kitaev and Klein, 2018). Although recent works (Lou et al., 2022; Yang and Tu, 2021) model nested NER with the similar idea, they ignore the global information in the corpus. Thus, we extend our idea of holistic structure modeling from sentence-level to corpus-level by introducing corpus-aware features according to the statistics of point-wise mutual information or frequency over dataset and aggregating the information to our model by span attention (Tian et al., 2020). The corpus-aware features hopefully enhance the current model by making up for such inability of the current representations.

Our experiments are conducted on three widely-used benchmark datasets following previous studies (Luo and Zhao, 2020; Wang et al., 2020). Results show that our model approaches or exceeds the current state-of-the-art with the newly-introduced methods. Our contribution can be summarized into three-fold:

- We take a holistic structure to model all nested named entities in a sentence and adopt a chart parsing algorithm to decode every levels of nested named entities once for all, avoiding the error propagation problem while offering a concise and natural solution for so compli-

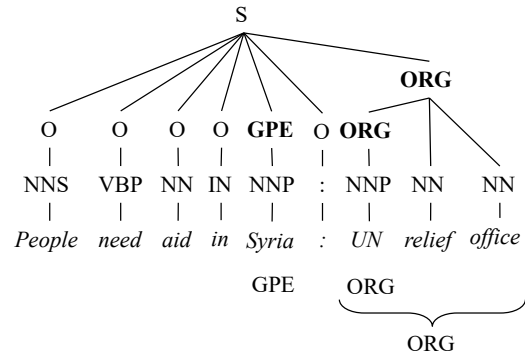


Figure 2: An example of nested NE tree. We use the tag ‘O’ to label the span that is not recognized as a named entity. The nested named entities form is transformed by removing the tag ‘O’.

cated a task.

- We introduce corpus-aware features derived from PMI or frequency statistics to further boost model performance and enable the model a good domain adaptation capability.
- With the simplicity of our modeling idea, our proposed holistic structure parsing model yields performance on par with or surpassing the state-of-the-art on widely-used English benchmark datasets.

2 Model

2.1 Holistic Structure Modeling for Nested Named Entities

Given a sentence of n words $\mathcal{X} = \{x_1, x_2, \dots, x_n\}$, we denote each named entity as a triplet (i, j, l) , where i, j are the beginning and ending positions of a NE with label $l \in \mathcal{L}$, and \mathcal{L} represents the set of NE types. Taking every levels of nested NE as constituent as syntactic parsing, all the labeled NE spans can form a tree structure like syntactic constituent parse tree (shown in Figure 2).

Our model assigns a real value score $s(i, j, l)$ to each labeled span. Then the score of a candidate tree \mathcal{T} can be computed by the scores from all spans inside the tree:

$$s(\mathcal{T}) = \sum_{(i,j,l) \in \mathcal{T}} s(i, j, l). \quad (1)$$

Our model is trained and predicts the structure by selecting the tree $\hat{\mathcal{T}}$ with the highest score by:

$$\hat{\mathcal{T}} = \arg \max_{\mathcal{T}} s(\mathcal{T}). \quad (2)$$

Following (Gaddy et al., 2018; Taskar et al., 2005), we apply margin-based training for such a structured prediction problem. Our model is trained to satisfy the constraint:

$$s(\mathcal{T}^*) \geq s(\mathcal{T}) + \Delta(\mathcal{T}^*, \mathcal{T}), \quad (3)$$

where \mathcal{T}^* is the golden tree and \mathcal{T} covers all valid trees. Δ is the Hamming loss on labeled spans. Such a constraint guarantees two points: (1) The score of the golden tree is always greater than or equal to the candidate trees. (2) It enforces the model to make the tree structures with high scores close to the golden structure. And our objective function can be defined by hinge loss:

$$\max \left(0, \max_{\mathcal{T}} [s(\mathcal{T}) + \Delta(\mathcal{T}, \mathcal{T}^*)] - s(\mathcal{T}^*) \right).$$

2.2 Contextualized Encoding

Our model is implemented in an encoder-decoder framework. For the encoder, it encodes each word x_t in the input sentence $\mathcal{X} = \{x_1, x_2, \dots, x_n\}$ into vectors h_t . Following (Fried et al., 2017; Kitaev and Klein, 2018), \mathcal{X} will be transformed into word embedding $\mathcal{W} = \{w_1, w_2, \dots, w_n\}$ by embedding model. In addition, we embed part-of-speech (POS) tag and word position in sentence as m_i and p_i respectively for x_i . We set these three embeddings to the same dimension and then add them together to generate the input vector $z_i = w_i + m_i + p_i$ for the encoder.

We adopt the Transformer (Vaswani et al., 2017) as our encoder which is composed of 8 modules of the same self-attention structure. Each structure contains a multi-head attention SubLayer₁ and a feed-forward SubLayer₂. After each sublayer, there will be a residual connection and a Layer Normalization. So when given an input x , the output y_2 of each module is:

$$y_2 = \text{LayerNorm}(y_1 + \text{SubLayer}_2(y_1)), \quad (4)$$

where LayerNorm denotes the Layer Normalization and y_1 is the output of the multi-head attention:

$$y_1 = \text{LayerNorm}(x + \text{SubLayer}_1(x)). \quad (5)$$

After 8 stacked modules, the final output h_t for word at position t is generated.

2.3 Decoder: Holistic Structure Parsing

Our decoder performs the needed holistic structure parsing for every levels of nested NEs in a sentence.

After obtaining the context-aware representation h_t for each word, the representation of NE span (i, j) is computed by:

$$r_{i,j} = h_j - h_i, \quad (6)$$

which assumes that the information of a subsequent step is generated by merging a previous state with the information of the span between them, so the difference between two steps can represent the span. In general, $s(i, j, \cdot)$ for span (i, j) over label set \mathcal{L} is calculated by feeding $r_{i,j}$ into a multi-layer perceptrons (MLP), with ReLU activated function, which can be formalized by:

$$o_{i,j} = \text{ReLU}(\text{LayerNorm}(W_1 \cdot r_{i,j} + b_1)), \quad (7)$$

Finally, we have:

$$s(i, j, \cdot) = W_2 \cdot o_{i,j} + b_2, \quad (8)$$

where W_1 , W_2 and b_1 , b_2 are all trainable parameters.

We adopt CKY parsing algorithm (Younger, 1967) as our decoding algorithm to search for the highest score $s(i, j)$ for span (i, j) . In details, to begin with the span $(i, i + 1)$ of length 1, we only need to consider its label:

$$s^*(i, i + 1) = \max_{l \in \mathcal{L}} s(i, i + 1, l). \quad (9)$$

Then, we can extract the highest score $s^*(i, j)$ for longer span (i, j) in a recursive way by searching for the best matching label l and the best boundary k :

$$s^*(i, j) = \max_{l \in \mathcal{L}} s(i, j, l) + \max_{i < k < j} [s^*(i, k) + s^*(k, j)]. \quad (10)$$

In this way, for the current span, the best label and best split point are chosen separately to find the highest score.

According to the above steps, for an entire sentence, we use a bottom-up approach to find the highest score $s^*(0, n)$ which is the sum of the highest scores for its subtrees. The full structure for the entire sentence can be rehabilitated by traversing backpointers. Thus we parse the sentence as a holistic structure and disclose all nested named entities by giving the nested NE tree once for all.

As the same as classical chart parsing, the computational complexity for the CKY algorithm to parse a sentence of length n is $O(n^3)$, which is better than enumerating all possible spans in $O(n^4)$.

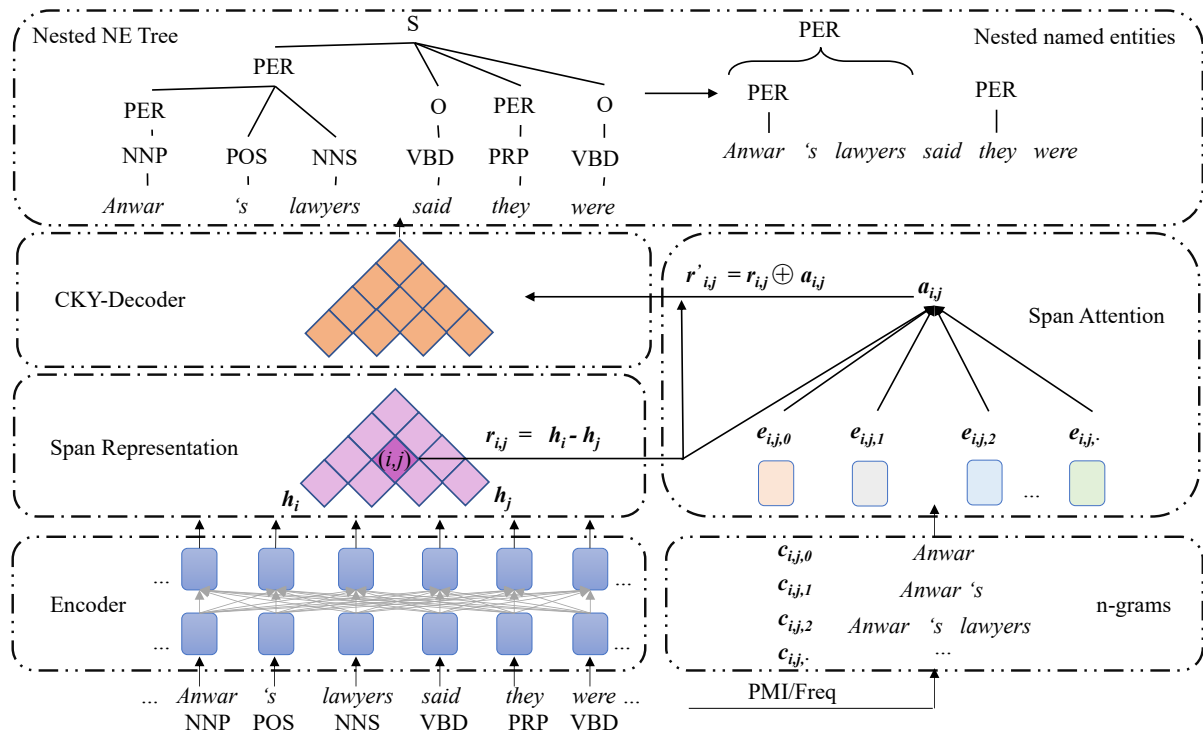


Figure 3: The architecture of our model, where the encoder converts the input sentence into a set of hidden states and spans are represented based on them. Then the span representation and the span attention calculated by n -gram are concatenated together and entered into CKY-decoder.

3 Corpus-aware Features

Though our proposed holistic structure parsing model at sentence-level hopefully captures the entire sentence for predicting all nested NEs, which is supposed to yield more powerful model capability than previous models, it still has the limitation similar with the previous ones which only learn the representation inside a sentence. Meanwhile, NEs quite sparsely distribute in all text data such as most NEs only occur for once in the entire corpus. Besides, there may be a huge domain difference among various corpora, which lets the model recognize new NEs even more difficultly. Thus, we further introduce corpus-aware clues to alleviate the above mentioned difficulties.

In detail, we consider corpus-level statistics including frequency and point-wise mutual information (PMI) over words and corresponding POS tags. Therefore, we consider the following statistics to form enhanced features, PMI on word (Word PMI), word frequency (Word Freq), PMI on POS tags (POS PMI) and on POS tags frequency (POS Freq).

Equipped with the corpus-aware features from a large amount of unlabeled corpus, the model may be enabled to effectively adapt to different domains by alleviating the difficulty caused by the sparsity

of NEs distribution.

3.1 Point-wise Mutual Information and Frequency

The PMI between two adjacent tokens x', x'' is calculated as:

$$\text{PMI}(x', x'') = \log \frac{p(x'x'')}{p(x')p(x'')} \quad (11)$$

We calculate the PMI between two adjacent words in turn and split a sentence from the low PMI positions to obtain multiple n -grams.

Similar with the above Word PMI processing, we can also count every n -grams ($n < 10$) to record their frequency in the corpus, then perform sentence segmentation according to a pre-specified threshold t_{freq} following the same processing to obtain the results of Word Freq.

In view of the sparsity of n -gram occurrence, n -gram statistics like the above Word PMI or Word Freq may still encounter the issue of data sparsity and long-tail effect (Kordumova et al., 2016). Thus we turn to the more informative POS which usually has a limited sized POS tag set (Benajiba and Rosso, 2007; Diab, 2007).

To apply the POS information, we first perform POS tagging over every sentences in the corpus

and save the correspondence between the word and the POS tag. Then we segment the spans of POS tags according to PMI or frequency. Finally, we convert these spans back to their corresponding word n -grams.

3.2 Integration of Corpus-aware Features

Following (Tian et al., 2020), we incorporate corpus-aware features into our model through span attention scoring to update the original NE span score in Eq. (6).

Before training our model, we extract all the n -grams from the training set and development set by Word PMI, Word Freq, POS PMI, POS Freq and store them in a Lexicon \mathcal{N} . Given a sentence \mathcal{X} , we enumerate all the spans (i, j) in it and find the spans which are included by \mathcal{N} to generate a set $\mathcal{C}_{i,j} = \{c_{i,j,1}, c_{i,j,2}, \dots, c_{i,j,v}, \dots, c_{i,j,m_{i,j}}\}$, where $m_{i,j}$ is the total amount of the spans. By span attention, each n -gram $c_{i,j,v}$ in $\mathcal{C}_{i,j}$ will be given an attention score by:

$$a_{i,j,v} = \frac{\exp(\mathbf{r}_{i,j}^\top \cdot \mathbf{e}_{i,j,v})}{\sum_{v=1}^{m_{i,j}} \exp(\mathbf{r}_{i,j}^\top \cdot \mathbf{e}_{i,j,v})}, \quad (12)$$

where $\mathbf{e}_{i,j,v} \in \mathbb{R}^{d_r}$ refers to the embedding of $c_{i,j,v}$. By the weighted average of n -gram embeddings, the resulted attention of span (i, j) is:

$$\mathbf{a}_{i,j} = \sum_{v=1}^{m_{i,j}} a_{i,j,v} \mathbf{e}_{i,j,v}. \quad (13)$$

As there are more short n -grams in the corpus than long n -grams, it is necessary to encourage these n -grams of various lengths in a balanced way, since there is no difference in the probability of being an entity. Thus, we split n -grams by their lengths into different categories, i.e., $\mathcal{C}_{i,j} = \{\mathcal{C}_{i,j,1}, \mathcal{C}_{i,j,2}, \dots, \mathcal{C}_{i,j,u}, \dots, \mathcal{C}_{i,j,n}\}$, where $u \in [1, n]$ indicates the n -gram length and weight them by each category following Eq. (12-13) to calculate $a_{i,j,v}^{(u)}$ and $\mathbf{a}_{i,j}^{(u)}$. Thus the final attention of the split n -grams is composed of the cascade all category attentions:

$$\mathbf{a}_{i,j} = \bigoplus_{1 \leq u \leq n} \delta_u \mathbf{a}_{i,j}^{(u)}, \quad (14)$$

where the weight of attentions from different categories is counterpoised by trainable parameter δ_u . After $\mathbf{a}_{i,j}$ is calculated for span (i, j) , our model concatenates it with $r_{i,j}$ (generated by Eq. 6): $\mathbf{r}'_{i,j} = \mathbf{r}_{i,j} \oplus \mathbf{a}_{i,j} \in \mathbb{R}^{d_r \cdot (n+1)}$. Then $\mathbf{r}'_{i,j}$ will be used to get $s(i, j, \cdot)$ by Eq. (7-8).

4 Experiments

4.1 Setup

We evaluate our approach on three nested named entity recognition benchmark datasets: GENIA, ACE2005 and KBP2017 datasets.

ACE2005 (Walker et al., 2006) contains 25% nested named entities and 7 entity types. We apply the same setup as (Katiyar and Cardie, 2018; Lu and Roth, 2015; Wang and Lu, 2018) by splitting the dataset into training/development/test sets by 8:1:1, respectively.

GENIA dataset (Kim et al., 2003) is based on the GENIAcorpus3.02p². The dataset contains 10% nested mentions and 5 entity types. We follow the same train/dev/test as previous work (Finkel and Manning, 2009; Lu and Roth, 2015) and split first 81% as training set, subsequent 9% as development set, and last 10% as test set.

KBP2017 contains 19% nested entities. We evaluate our model on the Event Nugget Detection Evaluation dataset (LDC2017E55) following (Lin et al., 2019), and previous annotated datasets (LDC2015E29, LDC2015E68, LDC2016E31 and LDC2017E02) are added into training and development sets. We split the datasets into 866/20/167 documents for training, development and test.

We use Precision (P), recall (R) and F-score (F₁) to evaluate the predicted named entities and we consider a named entity to be correctly recognized only if its boundary and category are both recognized correctly.

4.2 Implementation

In our experiments, we use randomly initialized embedding, variants of BERT (Devlin et al., 2019) and XLNet_{large} (Yang et al., 2019) to embed sentences. Following (Tian et al., 2020), we initialize all n -gram embeddings in the span attention module and match their dimensions with the hidden vectors from the encoder. During the training process, we use Adam optimizer with the learning rate 5e-5, 1e-5, and 5e-6. We also fine-tune dropout to attention, POS tag embedding and residual with the rate of 0.2, 0.4, 0.5. We select the model with the highest F₁ on the development set and evaluate it on the test set.

We set t_{pmi} to 0 to determine which words to form n -grams and 0.5 for POS tags since some POS tags often appear next to each other, but their

²<http://www.geniaproject.org/genia-corpus/posannotation>

Model	ACE 2005			GENIA			KBP 2017		
	P	R	F ₁	P	R	F ₁	P	R	F ₁
Hyper-Graph (Katiyar and Cardie, 2018)	70.60	70.40	70.50	77.70	71.80	74.60	–	–	–
Seg-Graph (Wang and Lu, 2018)	76.80	72.30	74.50	–	–	–	–	–	–
ARN (Lin et al., 2019)	76.20	73.60	74.90	75.80	73.90	74.80	77.70	71.80	74.60
Merge-BERT (Fisher and Vlachos, 2019)	82.70	82.10	82.40	–	–	–	–	–	–
DYGIE (Luan et al., 2019)	–	–	82.90	–	–	76.20	–	–	–
Path-BERT (Shibuya and Hovy, 2020)	82.98	82.42	82.70	78.07	76.45	77.25	–	–	–
BERT-MRC (Li et al., 2020)	87.16	86.59	86.88	85.18	81.12	83.75	82.33	77.61	80.97
Seq2seq-BART (Yan et al., 2021)	83.16	86.38	84.74	78.57	79.30	78.93	–	–	–
Seq2set (Tan et al.)	87.48	86.63	87.00	82.31	78.66	80.40	84.91	83.04	83.90
LexicalizedParsing (Lou et al., 2022)	85.97	87.87	86.91	78.39	78.50	78.44	–	–	–
Triaffine-ALBERT (Yuan et al., 2021)	88.88	88.24	88.56	–	–	–	89.42	85.22	87.27
Ours	86.81	88.70	87.75	79.76	75.74	77.70	88.25	87.10	87.67
-Corpus-Feat.	84.65	87.19	85.90	79.74	74.13	76.83	86.62	86.61	86.62

Table 1: Results for nested NER tasks.

DATA	Pre-trained	Word PMI			Word Freq.			POS PMI			POS Freq.			-Corpus-Feat.		
		P	R	F ₁	P	R	F ₁	P	R	F ₁	P	R	F ₁	P	R	F ₁
ACE2005	Random Init.	70.10	48.68	57.46*	73.35	46.10	56.62	72.13	49.31	58.58 †	74.33	48.22	56.53	68.35	47.21	55.85
	BERT _{base}	82.27	83.15	82.71 *	80.94	83.05	81.98	81.47	82.61	82.04†	81.25	82.31	81.78	78.06	79.34	78.69
	BERT _{warm}	84.05	85.47	84.75 *	83.42	85.29	84.34	83.42	85.29	84.34†	82.75	84.22	83.48	82.38	83.93	83.15
	XLNet _{large}	86.81	88.70	87.75 *	84.89	87.16	86.01	85.23	87.19	86.20†	84.95	87.40	86.16	84.65	87.19	85.90
GENIA	Random Init.	64.21	31.46	42.23*	71.05	29.44	41.63	61.85	43.97	51.40	73.03	47.28	57.40 †	60.05	28.29	38.46
	BERT _{base}	78.71	73.44	75.99 *	77.60	72.43	74.92	77.75	73.57	75.60†	77.67	71.60	74.51	76.86	69.38	72.93
	BERT _{warm}	79.07	74.92	76.94 *	78.79	74.26	76.46	78.38	74.76	76.53†	77.05	74.56	75.79	76.11	74.01	75.05
	XLNet _{large}	79.84	75.43	77.57	79.82	75.66	77.69*	79.76	75.74	77.70 †	79.21	75.30	77.21	79.74	74.13	76.83
KBP2017	Random Init.	75.73	49.32	59.73*	76.14	40.76	53.10	72.41	45.56	55.93	77.74	50.54	61.26 †	70.21	40.32	51.22
	BERT _{base}	83.45	80.22	81.80	82.95	81.18	82.06 *	82.42	80.10	81.24†	82.46	78.76	80.57	76.34	74.66	75.49
	BERT _{warm}	85.04	83.31	84.17	85.80	84.00	84.89 *	85.56	82.83	84.17†	82.91	79.35	81.09	80.45	79.48	79.96
	XLNet _{large}	88.15	86.48	87.30*	87.96	86.41	87.18	88.25	87.10	87.67 †	87.68	86.32	87.00	86.62	86.61	86.62

Table 2: Ablation study on the multiple corpus-aware features inducing approaches. * and † indicate the higher F₁ in Word PMI, Word Freq and POS PMI, POS Freq. Bold numbers indicate the highest F₁ in an experiment.

corresponding words are meaningless. Such as the POS tag pair (, VBD) has a high frequency of co-occurrence but the corresponding span (, had) is pointless, thus we set a higher threshold to filter this situation. Similarly, we set t_{freq} to 2 for words and 5 for POS tags.

4.3 Main Results

Table 1 compares our results with other state-of-the-art approaches on three nested NER benchmark datasets. By adding corpus-aware feature, our model has a significant improvement on KBP2017 which exceeds the state-of-the-art model (Yuan et al., 2021) by 1.88 R and 0.4 F₁. We also receive a comparable advancement on ACE2005 and outperform Yuan et al. (2021) by 0.46 R.

On the GENIA, our model performs relatively unsatisfactorily. But it is worth mentioning that the GENIA has less nested structure and is more close to flat NER, which indeed suits seq2seq (Tan et al.) modeling. And the rare biological nouns in the GENIA naturally fit for MRC method (Li et al.,

2020) that can better utilize the prior knowledge of pre-trained language model. Apart from these two strong baselines, our method still outperforms better than most models on the GENIA.

After removing the corpus-aware features of each model while the other training parameters remaining unchanged, the results show that the F₁ of all models receive a drop greatly, which proves that our corpus-aware features indeed enhance the effect. Although without them, our method also surpasses the previous best results on the KBP2017 by a large margin and surmounts most preceding approaches on ACE2005 and GENIA, which confirms modeling nested NER as holistic structure is essentially feasible.

5 Ablation Study

5.1 Effect of different corpus-aware features

For exploring the effects of the four corpus-aware features with significance test, we carry out twelve sets of experiments. For each set of them, we combine different features and pre-trained language

models to conduct comparative analysis.

The results are shown in Table 2. For each dataset, we use a randomly initialized word embedding and three pre-training language models to experiment separately and each pre-training language model has been trained with four corpus-aware features. During the training process, we allow the model to fine-tune the pre-trained language model in order to get better result. In order to better verify the effect of them, we also conduct a set of controlled experiment ablating corpus-aware features as a comparison.

In the comparison between PMI and Freq, the PMI-based features devote 8 highest F_1 in all 12 experiments while the Freq-based features occupy the remaining 4 highest results. Namely, PMI-based features show more helpful than Freq-based features for both Word and POS types, which implies PMI is indeed a better statistical measure than the common Freq for building informative corpus-aware features. The result also shows that text-level co-occurrence knowledge has more information than vanilla frequency.

5.2 Role of POS Tags

The comparison of Word and POS is shown in Figure 3. The Word-based features attain 7 highest F_1 while the POS-based features gain the rest 5 highest results. This phenomenon displays that Word-based features are slightly better than POS-based features. This conclusion is also in line with common perception that words themselves provide more knowledge than their POS tags.

The performance obtained by removing corpus-aware features is always the calamity which likewise supports the significant improvement of our proposed corpus-aware features. From overall view, the Word PMI has the best lifting effect for our model, followed by POS PMI. Word Freq and POS Freq are slightly less effective while both of them share the similar effect. Meanwhile, the proposed corpus-aware features can be easily imposed into any models that require span extraction.

5.3 Domain Adaptation Effects

In order to verify that the corpus-aware features can enhance the domain adaptability of our model, we conduct cross-domain experiments between the three datasets. We find the best trained model on one dataset and evaluate the effect of this model on other datasets as our baseline (shown at the top in Table 4). Subsequently, we retrain the best trained

Method	P	R	F_1
Ours	88.25	87.10	87.67
-POS embedding	87.64	86.21	86.92
-POS PMI	86.62	86.61	86.62
-Both	86.31	85.89	86.10

Table 3: Ablation study on POS tags.

model without changing any hyper-parameter setting and add external unlabeled data separately. In details, the model extracts n -gram from the external unlabeled data by leveraging the corpus-aware features and blends them into the lexicon N gleaned from the labeled training dataset. The changed lexicon N will be used to generate span attention and the domain adaptation ability of our model will be affected subsequently.

The lower part of Table 4 shows the results. It can be inferred that when a best trained model is joined a non-NER labeled dataset, the prediction effect of the model on this dataset will improve. The biggest improvement occurs when the KBP-trained and ACE-trained model are added GENIA non-NER labeled data. We deduce that the huge increments are because ACE and KBP both derive from the event domain while GENIA comes from a different biological domain. The model will perform poorly in a domain which is pretty different from the training corpus. Once the knowledge of that domain is compensated through the corpus-aware feature, the performance of the model in that domain will be greatly improved. This result verifies the effectiveness of our corpus-aware features for better domain adaptation.

We unexpectedly discover that the domain adaptation between ACE and KBP datasets has not made effective progress because they share the same domain. The corpus-aware features can not bring additional useful knowledge when the external unlabeled datasets are in the same domain with the training corpora. This inference is in line with our intuition and likewise demonstrates that our proposed corpus-aware features are able to extract information from different domains. We also find a counterintuitive fact that GENIA-trained model with KBP unlabeled data outperforms the GENIA-trained model with ACE unlabeled data on ACE. So we infer that the knowledge in the news domain brought by KBP unlabeled data suits well for this GENIA trained model and promotes its performance on another news domain ACE dataset.

In summary, the corpus-aware features can not

Model	ACE 2005			GENIA			KBP 2017		
	P	R	F ₁	P	R	F ₁	P	R	F ₁
ACE-trained Model	86.81	88.70	87.75	24.28*	10.19*	14.36*	73.89	76.06	74.96
GENIA-trained Model	6.49*	0.10*	0.19*	79.76	75.74	77.70	51.66*	2.18*	4.18*
KBP-trained Model	73.49	71.78	72.62	5.88*	0.88*	1.54*	88.25	87.10	87.67
<i>ACE-trained Model</i>									
+ GENIA Unlabeled Data	86.51	87.96	87.23	27.36*	14.09*	18.60*	74.40	74.60	74.50
+ KBP Unlabeled Data	85.72	87.50	86.60	22.50*	7.91*	11.71*	74.66	75.40	75.03*
<i>GENIA-trained Model</i>									
+ ACE Unlabeled Data	50.00*	1.20*	2.35*	79.38	74.80	77.02	37.64*	0.81*	1.59*
+ KBP Unlabeled Data	51.05*	2.53*	4.82*	78.13	74.01	76.01	52.02*	3.13*	5.90*
<i>KBP-trained Model</i>									
+ ACE Unlabeled Data	73.87	71.85	72.85	16.60*	3.86*	6.26*	87.28	86.86	87.07
+ GENIA Unlabeled Data	74.19	70.91	72.51	17.01*	4.28*	6.83*	86.46	84.79	85.61

Table 4: Results for domain adaptation experiments. * indicates unlabeled F₁ due to different named entity annotation standards. Bold numbers represent an improvement compared to the baselines.

only enhance the effect of the model but also can effectively improve its domain adaptation ability.

6 Related Work

- **Hypergraph model.** This approach is a common method applied for the nested named entities whose main idea is to apply hypergraph structure to represent nested NEs. Lu and Roth (2015) first proposed a hypergraph-based method by connecting multiple nodes with edges to represent nested NEs. Muis and Lu (2017) further employed multigraph representation and introduced a novel notion of mention separators to detect nested named entity. However, these methods are not intelligent enough due to the need to manually design explicit hypergraph.
- **Stacking layered model.** This is also a widespread approach to handle nested NER. Alex et al. (2007) stacked several neural layers to recognize the lower and higher level NEs separately. Fisher and Vlachos (2019) proposed a neural network that merges NEs or tokens to generate nested structure and labels them. Several recent researches applied multi-layer GCN to accomplish nested NER (Li et al., 2021; Luo and Zhao, 2020). These are practical methods, nonetheless, the existing models with a huge depth are computationally impractical.
- **Region-based model.** This method is another common method for nested NER, which enumerates all the possible subsequences and recognizes them. Sohrab and Miwa (2018) re-

garded all spans as potential NEs and used neural networks to recognize them. A layered model that enumerates all the potential subsequences while preserving the sequence structure was presented by (Wang et al., 2020). The biggest disadvantage of this method is that the model has computational complexity of at least $O(n^4)$ for the sentence of length n . Thus, it requires a lot of time overhead.

- **Parsing-based model.** Some concurrent work also model the nested NER as a parsing task. Lou et al. (2022) put forward a latent lexicalized parsing model for nested NER. Yang and Tu (2021) proposed a unified layer-by-layer bottom-up model for nested NER and constituency parsing. The formal additionally considers dependency syntax information, but its decoding speed is between $O(n^3)$ and $O(n^4)$, which is much lower than ours and not significantly better than our performance. The latter leverages the cascade decoding and suffers from the error propagation issue which does not exist in our model. Essentially different from us, none of them take into account the information of the corpus.

7 Conclusion

NER, so as to achieve the purpose of labeling all nested NEs once for all. On the basis of the proposed holistic sentence structure modeling, we further explore NER model enhancement from corpus-aware statistics with the hope of alleviating the serious sparsity issue of NEs. For this purpose, we extend our holistic structure modeling from sentence-level to corpus-level by offering multiple

corpus-aware features including Word Freq, Word PMI, POS Freq and POS PMI. The experimental results demonstrate the effectiveness of our proposed model by providing consistent and general performance improvement over strong baselines. In details, our model achieves performance approaching state-of-the-art for two datasets and reaches new state-of-the-art for one dataset. Last but not the least, our corpus-level holistic structure modeling shows surprising merit for effective domain adaptation.

References

- Beatrice Alex, Barry Haddow, and Claire Grover. 2007. [Recognising nested named entities in biomedical text](#). In *Biological, translational, and clinical language processing*, pages 65–72, Prague, Czech Republic. Association for Computational Linguistics.
- Yassine Benajiba and Paolo Rosso. 2007. Anersys 2.0: Conquering the ner task for the arabic language by combining the maximum entropy with pos-tag information. In *IICAL*, pages 1814–1823.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. [BERT: Pre-training of deep bidirectional transformers for language understanding](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 4171–4186, Minneapolis, Minnesota. Association for Computational Linguistics.
- Mona Diab. 2007. [Improved Arabic base phrase chunking with a new enriched POS tag set](#). In *Proceedings of the 2007 Workshop on Computational Approaches to Semitic Languages: Common Issues and Resources*, pages 89–96, Prague, Czech Republic. Association for Computational Linguistics.
- Jenny Rose Finkel and Christopher D. Manning. 2009. [Nested named entity recognition](#). In *Proceedings of the 2009 Conference on Empirical Methods in Natural Language Processing*, pages 141–150, Singapore. Association for Computational Linguistics.
- Joseph Fisher and Andreas Vlachos. 2019. [Merge and label: A novel neural network architecture for nested NER](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 5840–5850, Florence, Italy. Association for Computational Linguistics.
- Daniel Fried, Mitchell Stern, and Dan Klein. 2017. [Improving neural parsing by disentangling model combination and reranking effects](#). In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 161–166, Vancouver, Canada. Association for Computational Linguistics.
- David Gaddy, Mitchell Stern, and Dan Klein. 2018. [What’s going on in neural constituency parsers? an analysis](#). In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*, pages 999–1010, New Orleans, Louisiana. Association for Computational Linguistics.
- Zhiheng Huang, Wei Xu, and Kai Yu. 2015. Bidirectional lstm-crf models for sequence tagging. *arXiv preprint arXiv:1508.01991*.
- Meizhi Ju, Makoto Miwa, and Sophia Ananiadou. 2018. [A neural layered model for nested named entity recognition](#). In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*, pages 1446–1459, New Orleans, Louisiana. Association for Computational Linguistics.
- Arzoo Katiyar and Claire Cardie. 2018. [Nested named entity recognition revisited](#). In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*, pages 861–871, New Orleans, Louisiana. Association for Computational Linguistics.
- J-D Kim, Tomoko Ohta, Yuka Tateisi, and Jun’ichi Tsujii. 2003. Genia corpus—a semantically annotated corpus for bio-textmining. *Bioinformatics*, 19(suppl_1):i180–i182.
- Nikita Kitaev and Dan Klein. 2018. [Constituency parsing with a self-attentive encoder](#). In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 2676–2686, Melbourne, Australia. Association for Computational Linguistics.
- Svetlana Kordumova, Jan van Gemert, and Cees GM Snoek. 2016. Exploring the long tail of social media tags. In *International Conference on Multimedia Modeling*, pages 51–62. Springer.
- Sujan Kumar, Saha Sanjay, Chatterji Sandipan Dandapat, Sudeshna Sarkar, and Pabitra Mitra. 2008. A hybrid approach for named entity recognition in indian languages.
- Guillaume Lample, Miguel Ballesteros, Sandeep Subramanian, Kazuya Kawakami, and Chris Dyer. 2016. [Neural architectures for named entity recognition](#). In *Proceedings of the 2016 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 260–270, San Diego, California. Association for Computational Linguistics.
- Fei Li, ZhiChao Lin, Meishan Zhang, and Donghong Ji. 2021. A span-based model for joint overlapped and discontinuous named entity recognition. In *Proceedings of the 59th Annual Meeting of the Association for*

- Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 4814–4828.
- Xiaoya Li, Jingrong Feng, Yuxian Meng, Qinghong Han, Fei Wu, and Jiwei Li. 2020. [A unified MRC framework for named entity recognition](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 5849–5859, Online. Association for Computational Linguistics.
- Hongyu Lin, Yaojie Lu, Xianpei Han, and Le Sun. 2019. [Sequence-to-nuggets: Nested entity mention detection via anchor-region networks](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 5182–5192, Florence, Italy. Association for Computational Linguistics.
- Xiaohua Liu, Shaodian Zhang, Furu Wei, and Ming Zhou. 2011. Recognizing named entities in tweets. In *Proceedings of the 49th annual meeting of the association for computational linguistics: human language technologies*, pages 359–367.
- Chao Lou, Songlin Yang, and Kewei Tu. 2022. Nested named entity recognition as latent lexicalized constituency parsing. *arXiv preprint arXiv:2203.04665*.
- Wei Lu and Dan Roth. 2015. [Joint mention extraction and classification with mention hypergraphs](#). In *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing*, pages 857–867, Lisbon, Portugal. Association for Computational Linguistics.
- Yi Luan, Dave Wadden, Luheng He, Amy Shah, Mari Ostendorf, and Hannaneh Hajishirzi. 2019. [A general framework for information extraction using dynamic span graphs](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 3036–3046, Minneapolis, Minnesota. Association for Computational Linguistics.
- Ying Luo and Hai Zhao. 2020. [Bipartite flat-graph network for nested named entity recognition](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 6408–6418, Online. Association for Computational Linguistics.
- Xuezhe Ma and Eduard Hovy. 2016. [End-to-end sequence labeling via bi-directional LSTM-CNNs-CRF](#). In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1064–1074, Berlin, Germany. Association for Computational Linguistics.
- Aldrian Obaja Muis and Wei Lu. 2017. [Labeling gaps between words: Recognizing overlapping mentions with mention separators](#). In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, pages 2608–2618, Copenhagen, Denmark. Association for Computational Linguistics.
- Alan Ritter, Sam Clark, Oren Etzioni, et al. 2011. Named entity recognition in tweets: an experimental study. In *Proceedings of the 2011 conference on empirical methods in natural language processing*, pages 1524–1534.
- Takashi Shibuya and Eduard Hovy. 2020. [Nested named entity recognition via second-best sequence learning and decoding](#). *Transactions of the Association for Computational Linguistics*, 8:605–620.
- Mohammad Golam Sohrab and Makoto Miwa. 2018. [Deep exhaustive model for nested named entity recognition](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 2843–2849, Brussels, Belgium. Association for Computational Linguistics.
- Zeqi Tan, Yongliang Shen, Shuai Zhang, Weiming Lu, and Yueting Zhuang. [A sequence-to-set network for nested named entity recognition](#).
- Ben Taskar, Vassil Chatalbashev, Daphne Koller, and Carlos Guestrin. 2005. Learning structured prediction models: A large margin approach. In *Proceedings of the 22nd international conference on Machine learning*, pages 896–903.
- Yuanhe Tian, Yan Song, Fei Xia, and Tong Zhang. 2020. [Improving constituency parsing with span attention](#). In *Findings of the Association for Computational Linguistics: EMNLP 2020*, pages 1691–1703, Online. Association for Computational Linguistics.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Lukasz Kaiser, and Illia Polosukhin. 2017. [Attention is all you need](#). In *Advances in Neural Information Processing Systems 30: Annual Conference on Neural Information Processing Systems 2017, December 4-9, 2017, Long Beach, CA, USA*, pages 5998–6008.
- Christopher Walker, Stephanie Strassel, Julie Medero, and Kazuaki Maeda. 2006. Ace 2005 multilingual training corpus. *Linguistic Data Consortium, Philadelphia*, 57:45.
- Bailin Wang and Wei Lu. 2018. [Neural segmental hypergraphs for overlapping mention recognition](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 204–214, Brussels, Belgium. Association for Computational Linguistics.
- Jue Wang, Lidan Shou, Ke Chen, and Gang Chen. 2020. [Pyramid: A layered model for nested named entity recognition](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 5918–5928, Online. Association for Computational Linguistics.
- Hang Yan, Tao Gui, Junqi Dai, Qipeng Guo, Zheng Zhang, and Xipeng Qiu. 2021. [A unified generative framework for various ner subtasks](#). *arXiv preprint arXiv:2106.01223*.

Songlin Yang and Kewei Tu. 2021. Bottom-up constituency parsing and nested named entity recognition with pointer networks. *arXiv preprint arXiv:2110.05419*.

Zhilin Yang, Zihang Dai, Yiming Yang, Jaime G. Carbonell, Ruslan Salakhutdinov, and Quoc V. Le. 2019. [Xlnet: Generalized autoregressive pretraining for language understanding](#). In *Advances in Neural Information Processing Systems 32: Annual Conference on Neural Information Processing Systems 2019, NeurIPS 2019, December 8-14, 2019, Vancouver, BC, Canada*, pages 5754–5764.

Daniel H Younger. 1967. Recognition and parsing of context-free languages in time n^3 . *Information and control*, 10(2):189–208.

Zheng Yuan, Chuanqi Tan, Songfang Huang, and Fei Huang. 2021. Fusing heterogeneous factors with triaffine mechanism for nested named entity recognition. *arXiv preprint arXiv:2110.07480*.